Towards Development of Advanced Verification and Validation Procedures and Tools for the Certification of Learning Systems in Aerospace Applications

Stephen Jacklin, Johann Schumann, and Pramod Gupta
NASA Ames Research Center

Michael Richard, Kurt Guenther, and Fola Soares
Dryden Flight Research Center

The emergence of advanced autonomous aircraft and the proposed expansion of adaptive control capability have increased the desire to develop safe, reliable control systems for next generation aircraft and spacecraft. Under damage or failure conditions, future dynamically adaptive flight control systems maybe able to reassign control surface allocations in flight, or use integrated propulsion control to regain aircraft control. Adaptive controllers have been proposed to identify aircraft stability and control derivatives in-flight and also to recover aircraft control in the event of sudden control surface failure from hardware failure or damage.[refs] Advanced control algorithms to support fully autonomous aircraft and spacecraft operation have also been designed to allow vehicle self-health assessment, navigation, and mission planning to occur simultaneously and without human intervention.[refs] These adaptive control and diagnosis systems are generally referred to as learning systems because they monitor sensory feedback information to identify or learn new relationships between control inputs and the system state elements under control. Although these control methods afford the ability to compensate for unforeseen changes in the operating environment or vehicle state, they also introduce the possibility of adding unintended functionality as well.

A serious difficulty hindering the deployment of advanced, adaptive flight-critical software is the requirement to show that it can operate only as intended and with very high reliability under all operating conditions. Adaptive controllers that can make rapid and automatic adjustments to enable self-healing in the event of vehicle damage might also act to make a healthy aircraft un-flyable or a safety hazard to other vehicles. How can it be assured that safety-critical malfunctions can never occur? The software implementation must be thoroughly analyzed and checked to provide sufficient assurance of its intended functionality, safety, and the absence of unintended functionality.

The process of checking the correctness of software is termed verification and validation. As defined by the current standard for certification of digital flight software, RTCA DO-178B, verification is the evaluation of the results of a process to ensure correctness and consistency with respect to the inputs and standards provided to that process. Validation is defined as the process of determining that the requirements are the correct requirements and that they are

Infotech@Aerospace Conference
26-29 September 2005
Hyatt Regency Crystal City
Arlington, Virginia

complete. (RTCA (Radio Technical Commission for Aeronautics) is a private association of over 250 aeronautical organizations (established 1935) as a means of resolving aeronautical problems related to electronics and telecommunications. FAA Advisory Circular 20-115B specifies the use of DO-178B as one means of securing FAA certification of digital computer software).

With regard to software development, verification may be viewed as a process of testing the software at each stage of its development to make sure it has been programmed as specified in the software requirements document. Nevertheless, verification cannot be viewed as running test cases and comparing expected results to actual results because such testing can never reveal the absence of errors. Therefore, as suggested in DO-178B, verification objectives are satisfied through a combination of reviews, analyses, the development of test cases and procedures, and the subsequent execution of those test procedures.

Validation comprises the testing effort to assure that the verified software is able to accomplish the purpose as stated in the software requirements document. Validation failures are generally the result of the requirements being stated incorrectly or incompletely. Validation need not be done only at the end of the software development process, but also at intermediate levels in the software life cycle.

Currently, for dynamically adaptive, learning systems, there is no established way to meet the requirements expressed RTCA DO-178B. For this reason, NASA and other developers are seeking to create guidelines for the verification and validation of adaptive used in flight-critical applications.[Refs] The specific requirements for adaptive system certification, however, are not easily discernable. For conventional, non-adaptive software, the outputs are usually a purely deterministic function of the operating conditions and control inputs. Hence, verification and validation can be accomplished both by thoroughly testing the software with a parametric variation of the expected inputs and test conditions, together with a thorough review of software design and traceability to requirements analyses. Adaptive systems pose new challenges not only because they are non-deterministic systems, but also because the operating regimes under which they must operate may not be fully known, and hence, their behavior. The problem is basically one of ensuring adequate test coverage. For example, a learning system designed to use control surfaces in new ways if the aircraft becomes damages may seem to be a non-deterministic problem. Whereas the number of control surfaces that can fail is a finite set, damage to the wing or tail surfaces may involve aeroelastic considerations having an infinite number of possibilities. Nevertheless, such a system is still deterministic if it can be shown that for any defined state of damage the controller produces deterministic outputs. Therefore, adaptive systems require that new procedures and methods must be developed to show an alternate means of compliance to DO-178B.

The full paper will discuss the following topics:

First, a review of the current requirements posed by DO-178B will be presented with special emphasis given to those requirements that most pertain to the certification of adaptive systems. The objective will be to highlight those areas in which new procedures are most needed to facilitate adaptive system certification.

Second, the paper will discuss V&V problems specifically applicable to adaptive systems. This discussion will address the test coverage problem, algorithm learning stability, and identification convergence. An unfortunate situation that may arise is that the adaptive algorithm may over correct identified parameters or control actions, thereby causing large fluctuations in the control outputs. Instead of learning the dynamics of the new systems, the algorithm may not converge or jump erratically between values with little chance of recovery. Such a system is unstable. A related problem is that a learning algorithm may be so stable that adaptation occurs too slowly. The paper will discuss several ways to verify the stability of adaptive systems, including testing in simulation test beds and various means of analysis.

Last, the paper will present a brief overview of some advanced tools and software technologies currently being developed to radically reduce the time and cost of certifying safety-critical, learning systems. These technologies range from tools that find errors early in the design cycle through tools that find programming language errors and facilitate testing. These tools include a number of advanced formal software verification methods and new validation methods for expanding coverage or providing on-line information to assess the operation of the learning system in real time. This section will provides a look at some of these methods, but is by no means intended to be comprehensive of the work being done in this fast moving research area.